

10/11/00
JCESS U.S. PTO

REQUEST FOR FILING AN ORIGINAL UTILITY PATENT APPLICATION
UNDER 37 CFR 1.53(b)

Attorney Docket No.: 57983.000005
Client Reference No.: 12283RO

Inventor(s): Abel C. Dasylva; Delfin Y. Montuno; Yanhe Fan

Title: TECHNIQUE FOR ROUTING DATA WITHIN AN OPTICAL NETWORK

Assistant Commissioner
for Patents
Washington, D.C. 20231

Attn: Box PATENT
APPLICATION

Sir:

This is a request for filing an original utility patent application under 37 CFR 1.53(b). Attached are:

- [X] 57 pages of specification, claims, abstract.
- [X] 16 sheets of informal drawings.
- [X] Declaration & Power of Attorney Document.
- [] Verified Statement Claiming Small Entity Status - Small Business Concern.
- [] Verified Statement Claiming Small Entity Status - Independent Inventor.
- [X] Assignment Document along with Recordation of Assignment Form.
- [] Certified Copy of _____.
- [] Preliminary Amendment.
- [] Associate Power of Attorney.
- [X] Information Disclosure Statement, Form PTO-1449 and reference(s).
- [X] Return Receipt Postcard with check in the amount of \$_____

JCESS U.S. PTO
09/688090
10/11/00

Patent Application
Attorney Docket No.:
Client Reference No.:

The filing fee has been calculated as shown below:

	# of Claims		Extra Claims	Rate	Fee
Total Claims	24	Minus 20	4	x \$18.00	\$ 72.00
Independent Claims	2	Minus 3	0	x \$80.00	\$ 0.00
Basic Application Fee					\$ 710.00
If multiple dependent claims are present, add \$260.00					\$ 0.00
Subtotal					\$ 782.00
Subtract ¼ if Small Entity					\$ 0.00
Total Application Fee					\$ 782.00
Assignment Recordation Fee					\$ 40.00
TOTAL FEE DUE					\$ 822.00

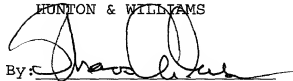
- [] Please charge Deposit Account No. 50-0206 in the amount of \$.00 for the above-indicated fees. A duplicate copy of this transmittal is submitted herewith.
- [X] The Commissioner is hereby authorized to charge any fees under 37 CFR 1.16 associated with the filing of this application, or credit any overpayment, to Deposit Account No. 50-0206. A duplicate copy of this transmittal is submitted herewith.
- [X] The Commissioner is hereby authorized to charge any fees under 37 CFR 1.16 and 1.17 which are required during the pendency of this application, or credit any overpayment, to

Patent Application
Attorney Docket No.:
Client Reference No.:

Deposit Account No. 50-206. This authorization does not include any issue fees under 37 CFR 1.18. A duplicate copy of this transmittal is submitted herewith.

Respectfully submitted,

HUNTON & WILLIAMS



By: Thomas E. Anderson
Registration No. 37,063

TEA/lhr

Hunton & Williams
1900 K Street, N.W.
Washington, DC 20006-1109
Telephone: 202-955-1500
Facsimile: 202-778-2201

Date: October 11, 2000

TECHNIQUE FOR ROUTING DATA WITHIN AN OPTICAL NETWORK

FIELD OF THE INVENTION

The present invention relates generally to optical
5 wavelength-switching and optical burst-switching and,
more particularly, to a label-switching technique for
routing data within an optical network.

BACKGROUND OF THE INVENTION

10 Emerging metropolitan optical networks are based on
optical cross-connects (OXC) controlled by internet
protocols (IPs). Both elements perform the critical task
of optical bandwidth provisioning. These networks
currently provide end-to-end circuit-switched bandwidth
15 allocation, through routing protocols like Open Shortest
Path First (OSPF), and through signaling protocols like
Multi-Protocol Label Switching (MPLS). The increasing
needs for greater network connectivity, and efficient
bandwidth utilization require more dynamic bandwidth
20 provisioning schemes such as optical burst-switching
(OBS) (see C. Qiao, and M. Yoo, "Choices, features and
issues in optical burst-switching", SPIE Optical Networks

Magazine, vol. 1, pages 36-44, April 2000). However,
this potential evolution does not change the need for
high-quality wavelength-switched optical service in the
near to medium term. In this context, it is appropriate
5 to design optical switching paradigms capable of
efficiently supporting both types of services. However,
the design of a common switch platform for wavelength-
switching and optical burst-switching (OBS) raises
problems, because of their quite different provisioning
10 time-scales. In wavelength-switching, lightpaths are
provisioned for hours, days or even months, while in
optical burst-switching (OBS) bandwidth is reserved on
burst time scales, which are measured in sub-milliseconds
units. State of the art optical cross-connects (OXC's)
15 built with micro-electro-mechanical systems (MEMS)
technology may have cross-connection latencies expressed
in tens of milliseconds (see N.R. Jankowski, C.
Bobrowski, D. Zipkin, R.R. Krchnavek, and R. Chamberlain,
"MEMS-based optical switch design for reconfigurable
20 fault-tolerant optical backplanes", Proceedings of the 6th
International Conference on Parallel Interconnects, pages

149-156, October 1999). These devices certainly have enough agility to support wavelength-switching, but cannot establish optical cross-connections on a per-burst basis as it is required by optical burst-switching (OBS).

5 In view of the foregoing, it would be desirable to provide optical wavelength-switching and optical burst-switching techniques for routing data within an optical network which overcome the above-described inadequacies and shortcomings.

10

SUMMARY OF THE INVENTION

According to the present invention, a technique for routing data within an optical network having a plurality of network nodes is provided. In one embodiment, the
15 technique is realized by receiving data at a first network node via a first optical signal having a first wavelength. The first wavelength corresponds to a first optical frequency, and the first optical frequency is mapped to a first binary representation. The first
20 binary representation is divided into a first plurality of fields, wherein at least one of the first plurality of

fields corresponds to a routing label in a first label
stack. A top routing label in the first label stack
indicates a second network node. Based at least
partially upon the top routing label, the data is
5 transmitted from the first network node to the second
network node via a second optical signal having a second
wavelength. The first wavelength may be either the same
as or different from the second wavelength.

In accordance with other aspects of the present
10 invention, the top routing label may be popped off the
first label stack so as to promote a next routing label
in the first label stack. If such is the case, the
second wavelength corresponds to a second optical
frequency, wherein the second optical frequency is mapped
15 to a second binary representation. The second binary
representation is divided into a second plurality of
fields, wherein at least one of the second plurality of
fields corresponds to a routing label in a second label
stack. The top routing label in the second label stack
20 indicates a third network node. Also, the top routing
label in the second label stack corresponds to the next

routing label in the first label stack. Further, if the
network accommodates 2^N frequencies in the form of,

$$f_i = f_0 + i \cdot \Delta f$$

wherein $i = 0, 1, \dots, 2^N - 1$, the second optical frequency may

5 be defined by,

$$f_{i_{out}} = f_0 + 2^L (f_{i_{in}} - f_0) \cdot 2^{N-L} \Delta f$$

and,

$$i_{out} = 2^L (i_{in} - 2^{N-L})$$

wherein $f_{i_{in}}$ represents the first optical frequency, i

10 represents the value of the top routing label in the

first label stack, and L represents the bit length of the
top routing label in the first label stack.

In accordance with further aspects of the present
invention, the top routing label in the first label stack
15 may be swapped with a new routing label when the first
label stack contains more than two routing labels. If
such is the case, the second wavelength corresponds to a
second optical frequency, wherein the second optical
frequency is mapped to a second binary representation.

20 The second binary representation is divided into a second
plurality of fields, wherein at least one of the second

plurality of fields corresponds to a routing label in a second label stack. A top routing label in the second label stack indicates a third network node. Also, the top routing label in the second label stack corresponds to the new routing label. Further, if the network accommodates 2^N frequencies in the form of,

$$f_i = f_0 + i \cdot \Delta f$$

wherein $i = 0, 1, \dots, 2^N - 1$, the second optical frequency may be defined by,

$$f_{i_{out}} = f_{i_{in}} + 2^{N-L} \left(-l^1 \right) \Delta f$$

and,

$$i_{out} = i_{in} + 2^{N-L} \left(-l^1 \right)$$

wherein $f_{i_{in}}$ represents the first optical frequency, l^1 represents the value of the top routing label in the first label stack, l represents the value of the new routing label, and L represents the bit length of the top routing label in the first label stack.

In accordance with still further aspects of the present invention, a new routing label may be pushed onto the first label stack. If such is the case, the second wavelength corresponds to a second optical frequency,

wherein the second optical frequency is mapped to a
second binary representation. The second binary
representation is divided into a second plurality of
fields, wherein at least one of the second plurality of
5 fields corresponds to a routing label in a second label
stack. A top routing label in the second label stack
indicates a third network node. Also, the top routing
label in the second label stack corresponds to the new
routing label. Further, if the network accommodates 2^N
10 frequencies in the form of,

$$f_i = f_0 + i \cdot \Delta f$$

wherein $i = 0, 1, \dots, 2^N - 1$, the second optical frequency is
defined by,

$$f_{i_{in}} = f_0 + \left\lfloor \frac{(i_{in} - f_0)}{\Delta f} \right\rfloor 2^{-L} \Delta f + 2^{N-L} L \Delta f$$

15 and,

$$i_{out} = \left\lfloor \frac{i_{in}}{2^L} \right\rfloor + 2^{N-L} L$$

wherein $f_{i_{in}}$ represents the first optical frequency, i
represents the value of the top routing label in the
second label stack, and L represents the bit length of
20 the top routing label in the second label stack.

00000000-101000

In accordance with still further aspects of the present invention, at least another one of the first plurality of fields may correspond to a termination field indicating an end of the first label stack. Also, at least another one of the first plurality of fields may correspond to a contention field for use in differentiating the first wavelength from a third wavelength. If such is the case, the data is defined as a first data, and a second data is received at the first network node via a third optical signal having the third wavelength, wherein the first optical signal and the third optical signal have similar routing paths through the network.

The present invention will now be described in more detail with reference to exemplary embodiments thereof as shown in the appended drawings. While the present invention is described below with reference to preferred embodiments, it should be understood that the present invention is not limited thereto. Those of ordinary skill in the art having access to the teachings herein will recognize additional implementations, modifications,

and embodiments, as well as other fields of use, which are within the scope of the present invention as disclosed and claimed herein, and with respect to which the present invention could be of significant utility.

5

BRIEF DESCRIPTION OF THE DRAWINGS

In order to facilitate a fuller understanding of the present invention, reference is now made to the appended drawings. These drawings should not be construed as
10 limiting the present invention, but are intended to be exemplary only.

Figure 1 shows a conventional label stack in a data packet.

Figure 2 shows a conventional label-swap operation.

15 Figure 3 shows a conventional label-pop operation.

Figure 4 shows a conventional label-push operation.

Figure 5 shows a system wherein label-switched routing is performed using a conventional label-swap operation.

20 Figure 6 shows a system wherein label-switched routing is performed using a conventional label-pop

operation.

Figure 7 shows a wavelength-switched WDM system wherein conventional wavelength-switching is performed.

Figure 8 illustrates the two-way reservation process that takes place in the wavelength-switched WDM system of Figure 7.

Figure 9 shows a burst-switched WDM system for illustrating a conventional burst-switched transmission.

Figure 10 shows a label stack format in accordance with the present invention.

Figure 11 shows a label-swap operation in accordance with the present invention.

Figure 12 shows a label-pop operation in accordance with the present invention.

Figure 13 shows a label-push operation in accordance with the present invention.

Figure 14 shows an example of wavelength contention resolution in accordance with the present invention.

Figure 15 shows a wavelength-switched WDM system that employs a label-pop operation in accordance with the present invention.

Figure 16 illustrates the two-way reservation process that takes place prior to the establishment of a connection in the system of Figure 15.

Figure 17 shows a wavelength-switched WDM system that employs a label-swap operation in accordance with the present invention.

Figure 18 illustrates the two-way reservation process that takes place prior to the establishment of a connection in the system of Figure 17.

Figure 19 shown a wavelength-switched WDM system that employs both the label-push operation and the label-pop operation to provide transparent wavelength-trunking between two metropolitan area networks.

Figure 20 illustrates the two-way reservation process that takes place prior to the establishment of a connection in the system of Figure 19.

Figure 21 shows an optical burst-switched WDM system that employs a label-pop operation to successfully switch an optical burst in accordance with the present invention.

Figure 22 shows the optical burst-switched WDM

system of Figure 21 wherein an optical burst is
unsuccessfully switched due to a wavelength already
having been assigned to an ongoing circuit.

5 **DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENT(S)**

Label-switching comprises of all the data routing
and switching paradigms where packets are forwarded to
next-hop nodes according to specific labels. In these
schemes, there is no route computation on a per-packet
10 basis, and there is often a one-to-one mapping between
labels and outgoing interfaces at network nodes. Label-
switching enables the network to replace computation
intensive layer 3 operations (e.g. route computation) by
simpler layer 2 operations, which are implemented with
15 dedicated high performance hardware.

Label-switching relies on three fundamental
operations, which are label-swap, label-pop, and label-
push. To explain these operations, we assume that in
each packet header, the labels are organized into a
20 stack, as shown in Figure 1. That is, Figure 1 shows a
data packet 10 comprising a payload 12 and a header 14,

wherein the header 14 comprises a label stack 16 including a plurality of labels 18.

The label-swap operation simply replaces the top label 18a of the stack 16 with a new label 20 contained in a switching table (not shown), as shown in Figure 2. The binding between the old label 18a and the new label 20 is usually defined at connection setup.

The label-pop operation simply pops the top label 18a of the stack 16, as shown in Figure 3.

The label-push operation simply inserts a new label 24 into the stack 16, as shown in Figure 4.

The joint use of the label-pop operation and the label-push operation enables the definition of virtual paths.

The definition of routing labels may be static or dynamic. In the static case, the labels and the mapping between labels and next-hop nodes are based on configuration information, as in the case of MPLS hop-by-hop routing. When labels are defined dynamically, the labels or the binding between labels and next-hop nodes are set on a per-flow basis, by a signaling protocol.

Figure 5 shows a system 30 wherein label-switched routing is performed using the label-swap operation.

That is, the system 30 comprises a source 32, a destination 34, and a network 36, which includes a first edge node 38, a second edge node 40, and a plurality of routers 42 interconnected by a plurality of interconnects 44. As a data packet advances across the network 36 from the source 32 to the destination 34, the top label of the label stack is simply replaced with a new label contained in a switching table, as shown (e.g., at router D, the top label 11 from incoming router B is replaced with new top label 01 for outgoing router E). As mentioned above, the binding between the old label and the new label is usually defined at connection setup.

Figure 6 shows a system 50 wherein label-switched routing is performed using the label-pop operation. The system 50 is physically identical to the system 30 and thus is numerically identified as such. However, in the system 50, as a data packet advances across the network 36 from the source 32 to the destination 34, the top label of the label stack is simply popped off the label

stack, as shown (e.g., at router D, the top label 11, indicating router E, is popped off the label stack). This operation reduces the header overhead of packets as they get closer to egress nodes.

5 The label-switching paradigm has also been extended to wavelength division multiplexing (WDM) networks. In fact, IP label-switching protocols have been extended to support wavelength-switching, where lightpaths are established and deleted according to a circuit-switching
10 paradigm. These extensions use wavelengths as labels, and have been called optical label-switching protocols (see N. Ghani, "Lambda-labeling: A framework for IP-over-WDM using MPLS", SPIE Optical Networks Magazine, vol. 1, pages 45-58, April 2000). However, these protocols do
15 not explicitly encode route information in wavelengths or interfaces (i.e., it is impossible to predict the next hop of an optical packet by simply considering its incoming interface and its wavelength). Instead, the assignment of wavelengths to outgoing interfaces is
20 assumed to be utterly arbitrary, and only decided at connection setup, on a per-flow basis. Although waveband

routing appears as an option of current optical label-switching protocols, these protocols do not address the appropriate definition of wavebands to provide source-routing for optical burst-switching (OBS) or circuit-switched wavelength routing.

Referring to Figure 7, there is shown a wavelength-switched WDM system 60 that may be operating under any of the protocols that have heretofore been proposed for optical label-switching. The wavelength-switched WDM system 60 comprises a source 62, a destination 64, and a network 66, which includes a first edge node 68, a second edge node 70, and a plurality of optical cross-connects (OXC) 72 interconnected by a plurality of interconnects 74. As shown in Figure 7, as a data packet advances across the network 66 from the source 62 to the destination 64, the wavelength at which the data packet is transmitted changes.

It is assumed that there is a request for a lightpath with a bandwidth of one wavelength between the source 62 and the destination 64. The assignment of wavelengths to the different portions of the lightpath is

only based on the availability of the wavelength channels, and on the wavelength-conversion capabilities.

Bindings between incoming and outgoing channels are stored in connection tables, as shown in Figure 7. The
5 connection table entries corresponding to the lightpath request between the source 62 and the destination 64 are represented by giving the incoming node and wavelength, as well as the outgoing node and wavelength. For example, at a first optical cross-connect 72 (i.e., OXC
10 B), the entry $(A, F_{11}) : (D, F_{11})$ means that the channel coming from the first edge node 68 at optical frequency F_{11} is to be switched to a second optical cross-connect 72 (i.e., OXC node D) at frequency F_{11} . Figure 8 illustrates the two-way reservation process that takes place prior to the
15 establishment of the connection between the source 62 and the destination 64 in the system 60 of Figure 7.

Optical burst-switching (OBS) is based on a tell-and-go transmission model wherein traffic sources send control packets ahead of data bursts and do not wait for
20 resource confirmation messages before transmitting data packets (see C. Qiao, and M. Yoo, "Choices, features and

issues in optical burst-switching", SPIE Optical Networks Magazine, vol. 1, pages 36-44, April 2000). The burst transmission latency is the amount of time that separates control packets and data bursts. It is computed at

5 network ingress nodes to give the network enough time to setup the requested circuits. In WDM networks, optical burst-switching (OBS) control packets request wavelength cross-connections on a per-burst basis. Assuming that the OXCs in the network have a cross-connection latency
10 (i.e., the time it takes to connect two physical ports in an OXC) of δ_x , a message processing delay of δ_m , and an H hop path with propagation delay of δ_p , the burst transmission latency must be at least equal to

$$\delta_{\min} = H(\delta_x + \delta_m) - \delta_p.$$

15 Referring to Figure 9, there is shown a burst-switched WDM system 80 for illustrating a burst-switched transmission. The system 80 is physically identical to the system 60 and thus is numerically identified as such.

Also, similar to the system 60, in the system 80, as a
20 data packet advances across the network 66 from the source 62 to the destination 64, the wavelength at which

the data packet is transmitted changes. However, in the system 80, the data packets are sent in bursts preceded by control packets.

The efficiency of the burst-switching paradigm decreases as the burst transmission latency increases. Unfortunately, state-of-the art OXCs still have cross-connection latencies (δ_x) in the tens of milliseconds range, which is a dominant parameter of the overall burst-transmission latency and imposes an upper limit on attainable throughputs.

To overcome the problems associated with the aforementioned conventional label-switching schemes, a new label-switching paradigm is disclosed. This new label-switching paradigm involves the explicit encoding of routing information through specific wavelength assignments. That is, in this new label-switching paradigm, a wavelength is equivalent to a binary string.

This binary string is partitioned into routing bits, and contention resolution bits. The two types of bits are handled by disjoint operations to meet the routing and the contention resolution needs of the network. Thus,

this new label-switching paradigm is realized via a new label-stack format, and the processing of optical labels through specific label-stack operations, which are both described in detail below.

5 At the outset, it is assumed that the new label-switching paradigm is operating in a dense wavelength division multiplexing (DWDM) network with $W=2^N$ optical frequencies of the form $f_i = f_0 + i\Delta f$, where $i=0, \dots, 2^N-1$. It is also assumed that this network has proper wavelength-
10 conversion facilities at all network nodes to support the different operations described below.

Referring to Figure 10, there is shown a label stack format 90 in accordance with the present invention. To encode the label stack, each optical frequency, f_i , is
15 mapped to the binary representation of i . For example, when $N=8$, frequency f_{161} is mapped to binary number 10100001. The binary number of the optical frequency is then divided into smaller fields corresponding to routing labels. The obtained set of labels (i.e., Label 1, Label
20 2, etc.), read from left to right, and an optional stack termination string (i.e., End label), define the label

stack of the wavelength channel. The lengths of the routing labels and their interpretation may be node and interface dependent. Therefore, the label stack associated with a particular wavelength depends on the
5 selected lightpath.

To illustrate the foregoing, first consider $N=8$, a fixed label length $L=2$, and a stack termination string equal to 00. Then, the label stack of frequency f_{161} includes two labels that are $l_{161}^1=10$, $l_{161}^2=10$. However, if
10 the lightpath goes through two nodes supporting label-switched routing, such that the first node uses labels of length $L_1=1$, and the second node uses labels of length $L_2=3$, two different labels are obtained for f_{161} , which are $l_{161}^1=1$, $l_{161}^2=010$.

15 The label-stack operations include label-swap, label-pop, and label-push operations. The label-swap operation replaces the first label of the stack with another label, to be used by the next-hop. This operation requires that the swapped labels have the same
20 length. The label-swap operation can be mathematically defined as follows.

Consider some frequency f_{i_n} , such that the first label of the stack is l^1 of length L , and some other label $0 \leq l < 2^L$. When l^1 is replaced with l , the frequency $f_{i_{out}}$ is obtained which satisfies the following equations:

5

$$f_{i_{out}} = f_{i_n} + 2^{N-L}(-l^1)M$$

$$i_{out} = i_n + 2^{N-L}(-l^1)$$

Figure 11 further illustrates the label-swap operation
10 100 wherein the first label (i.e., Label 1) in an initial label stack 102 is replaced with a new label (i.e., New label) in a resulting label stack 104.

To give an example of the label-swap operation,
consider, as before, $N=8$, a fixed label length $L=2$, and
15 a stack termination string equal to 00. The first label of the stack of frequency $f_{i_{61}}$ is $l^1_{61}=10$. When this first label is replaced (i.e., swapped) with a new label $l=01$, the output frequency $f_{i_{out}}=f_{97}$ is obtained.

In the label-pop operation, labels that have already
20 been used can be removed from the label stack without any

impact on the optical service received by a wavelength channel. The removal of such useless labels from the stack simplifies the task of subsequent network nodes. To remove labels from the stack, a label-popping

5 wavelength-conversion is performed.

The precise mathematical definition of the label-pop operation for a label, l , of length L (in number of bits) is given by the following equations:

10

$$f_{i_{out}} = f_0 + 2^L (f_{i_{in}} - f_0) - 2^{N-L} L \Delta f$$
$$i_{out} = 2^L (i_{in} - 2^{N-L} L)$$

In this definition, it is assumed that l is the first label of the label stack of f_{in} . Figure 12 further
15 illustrates the label-pop operation 110 wherein the first label (i.e., Label 1) in an initial label stack 112 is removed (i.e., popped) in a resulting label stack 114.

To give an example of the label-pop operation, consider, as before, $N=8$, a fixed label length $L=2$, and
20 a stack termination string equal to 00. When the frequency f_{161} is popped, the output frequency f_{132} is

obtained. It is important to note that the label-pop operation also affects the wavelength bits that are not part of the label stack. However, this operation can be defined in a way that does not affect the bits outside of

5 the label stack.

The label-push operation is opposite to the label-pop operation as it introduces new labels into the stack.

The label-push operation is useful to introduce new routing labels into the label stack at any point along

10 the lightpath. However, pushed labels must be defined through prior signaling or recovered from data packets by other means than wavelength encoding.

The precise mathematical definition of the label-push operation for a label l of length L (in number of

15 bits) is given by the following equations:

$$f_{out} = f_0 + \left\lfloor \frac{(f_{in} - f_0)}{\Delta f} \right\rfloor 2^{-L} \Delta f + 2^{N-L} L \Delta f$$

$$i_{out} = \left\lfloor \frac{i_{in}}{2^L} \right\rfloor + 2^{N-L} L$$

20 Figure 13 further illustrates the label-push operation

120 wherein a new label (i.e., Label 0) in a resulting label stack 122 is introduced (i.e., pushed) over the first label (i.e., Label 1) in an initial label stack 124.

5 To give an example of the label-push operation, consider, as before, $N=8$, a fixed label length $L=2$, and a stack termination string equal to 00. When the label $l=10$ is pushed into frequency f_{160} , the output frequency f_{168} is obtained.

10 In wavelength-switching networks, it is necessary to avoid wavelength contention. When such contention occurs, it can be resolved by changing the wavelength bits that are not meaningful from the point of view of routing, which are hereafter referred to as wavelength
15 contention resolution bits. For example, consider, as before, $N=8$, a fixed label length $L=2$, and a stack termination string equal to 00. Also, assume that two wavelength-channels share the same network path starting from some common downstream node, and are both assigned
20 the frequency f_{161} in the hop prior to a common node. To avoid collision on the next-hop link, one of these

channels may be assigned to another frequency (e.g., frequency f_{163}) before undergoing any label-stack operation. Of course, the newly assigned frequency (i.e., frequency f_{163}) must be available on the link. It is easy to check that the wavelength-interchange that occurs does not modify the original routing information contained in the frequency.

Referring to Figure 14, a representation 130 of the above-described wavelength contention resolution example is provided wherein data packets are being transmitted from both a first OXC 132 (i.e., OXC A) and a second OXC 134 (i.e., OXC B) at the frequency f_{161} to a third OXC 136 (i.e., OXC C). The third OXC 136 assigns the transmission from the first OXC 132 to another frequency (e.g., frequency f_{163}) before undergoing any label-stack operation to avoid collision on the next-hop link. In this particular example, the third OXC 136 assigns the new frequency by changing the two least significant wavelength bits (i.e., the wavelength contention resolution bits). The third OXC 136 then performs a label-pop operation on the incoming wavelength-channels

before forwarding them to a fourth OXC 138 (i.e., OXC D).

In the above-described example, it is assumed that the mapping between routing labels and next-hop nodes is the same for all incoming interfaces.

5 The above-described new label-switching paradigm may support wavelength-switching and wavelength-trunking, as well as optical burst-switching. For all of these label-switching applications, the new label-switching paradigm offers many options to establish lightpaths, depending on
10 the type of label-switching operation used at the different hops of a path, and on the nature of the network hops. That is, in some cases each hop may map to a physical link, while in other cases some hops may be logical and correspond to virtual paths. In the
15 following description, examples of each of the above-mentioned label-switching applications are provided, although it should be noted that the present invention is not limited to these specific examples.

Referring to Figure 15, there is shown a wavelength-
20 switched WDM system 150 that employs the label-pop operation in accordance with the present invention. The

wavelength-switched WDM system 150 comprises a source 152, a destination 154, and a network 156, which includes a first edge node 158, a second edge node 160, and a plurality of optical cross-connects (OXC) 162 interconnected by a plurality of interconnects 164. In the wavelength-switched WDM system 150, a number $W = 256$ (i.e., $N = 8$) of consecutive regularly spaced wavelength channels are assumed, as is a label length $L = 2$. Also, it is assumed that the traffic demand comprises one wavelength channel between the source 152 and the destination 154.

The source 152 and the destination 154 access the core network through the first edge node 158 and the second edge node 160, respectively. In the core network, the plurality of optical cross-connects (OXC) 162 provide waveband-routing by employing the label-pop operation in accordance with the present invention. The mapping between next-hop nodes and labels is given for each node 162 in the core network by the small table located next to each node 162 in the core network, as shown in Figure 15. For example, in the case of node OXC B, the mapping

is the following:

Next-hop A: 01

Next-hop C: 10

5 Next-hop D: 11

20 The correspondence between labels and wavebands is easily determined. For example, considering an arbitrary routing label b_1b_0 , the corresponding waveband is $[b_1b_0000000, b_1b_0111111]$. In other words, all frequencies in the band $[b_1b_0000000, b_1b_0111111]$ are sent to the next-hop of the routing label b_1b_0 , after being label-popped. Therefore, the mapping between next-hop nodes and routing labels can be specified by giving the mapping between next-hop nodes and wavebands. In the case of node OXC B, the equivalent description is the following:

Next-hop A: $[F_{64}, F_{127}]$

Next-hop C: $[F_{128}, F_{191}]$

20 Next-hop D: $[F_{192}, F_{255}]$

After receiving a connection request from the source 152, the first edge node 158 computes a path for the connection through the network 156. Assuming that this path is through OXC B, OXC D, OXC E, and OXC F, the first 5 edge node 158 sends a path setup request along the selected path. This path setup request contains a list of wavelengths which may be assigned at each node 162 along the path, taking into account the following constraints: 1.) At each hop, the assigned channel must 10 be available; 2.) For each node 162, the outgoing assigned wavelength channel is obtained from the incoming wavelength-channel through label-pop in the nodes 162 in the core network (i.e., it is not necessary to pop labels when the next-hop is either the first edge node 158 or 15 the second edge node 160, although it is possible to do so); and 3.) The path defined by assigned wavelength-channels must match the path selected by the first edge node 158. Figure 16 illustrates the two-way reservation process that takes place prior to the establishment of 20 the connection between the source 152 and the destination 154 in the system 150 of Figure 15.

The result of this path setup process depends on the current congestion in the network 156. For simplicity, it is assumed that the result of the path setup phase is the allocation of frequency F_{255} between OXC A and OXC B, 5 frequency F_{252} between OXC B and OXC D, frequency F_{240} between OXC D and OXC E, and frequency F_{240} between OXC E and OXC F.

The label stack encoded in frequency F_{255} can be checked to see if it matches the selected path. For 10 example, the binary representation of F_{255} is 11111111, which corresponds to four consecutive routing labels that are respectively $l_1=11$, $l_2=11$, $l_3=11$, and $l_4=11$. The first label $l_1=11$ is used at OXC B to select OXC D, as the label $l_1=11$ maps to OXC D for the next-hop. The application of 15 the label-pop operation on F_{255} yields F_{252} . Similar checks may be done for the other nodes 162 of the path.

Referring to Figure 17, there is shown a wavelength-switched WDM system 170 that employs the label-swap operation in accordance with the present invention. The 20 system 170 of Figure 17 is physically identical to the system 150 of Figure 15 and thus is numerically

identified as such. However, in the system 170 of Figure 17, the wavelength-switched lightpath is formed by employing the label-swap operation, rather than the label-pop operation of the system 150 of Figure 15.

5 In the system 170 of Figure 17, the input wavelength does not carry the whole route information. Instead, the binding between incoming and outgoing routing labels is stored in connection tables of the nodes 162 of the path at connection setup. At each hop of the path, the
10 wavelength of the lightpath is modified to carry the correct routing label for the next hop through the label-swap operation, and to avoid wavelength-collision with other ongoing circuits. Such a wavelength-interchange operation may be seen as a two-step process wherein the
15 following two operations occur in sequence: 1.) Modification of the wavelength bits carrying the routing labels; and 2.) Modification of the wavelength contention resolution bits to avoid collision. The two operations operate on disjoint sets of wavelength bits.

20 The mapping between next-hop nodes and labels is given for each node 162 in the core network by the small

table located next to each node 162 in the core network,
as shown in Figure 17. Also shown next to each node 162
in the core network is the connection table entry
corresponding to the selected lightpath between the
5 source 152 and the destination 154 when the node 162
belongs to the lightpath. As with the system 150, it is
easy to check that, at each hop of the lightpath, the
routing label of the wavelength is swapped according to
the target next hop node, and that no wavelength-
10 collision occurs. Figure 18 illustrates the two-way
reservation process that takes place prior to the
establishment of the connection between the source 152
and the destination 154 in the system 170 of Figure 17.

In both of the above-described cases of wavelength-
15 switched lightpaths (i.e., those utilizing the label-pop
operation and the label-swap operation), the route
computation algorithm generally must include the
constraints corresponding to the mapping between
wavebands and next-hop nodes which are imposed by the new
20 label-switching paradigm in accordance with the present
invention. In both cases, the route information is

advertised at connection setup, and the wavelength channels must be explicitly allocated by network nodes.

In addition to the above-described cases of wavelength-switched lightpaths (i.e., those utilizing the label-pop operation and the label-swap operation), wavelength-switched lightpaths may also be realized utilizing a combination of the label-push operation and the label-pop operation. This additional wavelength-switching application, termed wavelength-trunking, is particularly beneficial in local and metropolitan area networks. That is, the costs of WDM systems increase with the granularity of wavelength channels. In local and metropolitan area networks, the costs of WDM nodes are still a determining factor of the commercial use of WDM. To address the cost issue, coarse WDM has been proposed where channel spacings are much larger than in long-haul dense WDM systems. This difference in wavelength-channel granularity between the two types of networks reflect deeper functional differences. Local and metropolitan area networks put a stress on per wavelength routing and switching functions, while long-

haul networks have a more point-to-point nature and perform per wavelength operations only at the transmission layer to monitor the integrity of the optical signals. However, long-haul networks are

5 evolving to perform a certain amount of wavelength-routing and switching. The large numbers of wavelength-channels, which may be soon supported by long-haul systems, prohibit the same level of functionality on a per wavelength-basis as in local or metropolitan

10 networks. Therefore, waveband-switching schemes are appropriate to enable some wavelength-switching functionality in long-haul networks, with the definition of wavelength trunks corresponding to wavebands.

To implement such waveband-switching schemes, a

15 combination of the label-push operation and the label-pop operation is used where a label is pushed/popped at network ingress/egress nodes according to the waveband or virtual path to which it is assigned. As in traditional data networks, virtual paths are defined by traffic

20 engineering and traffic management activities. An important future revenue generating service of

metropolitan area networks is the definition of virtual private optical networks, where different nodes are spread over different metropolitan networks connected by long-haul trunks. For the implementation of such

5 services, wavelength-trunking based on label-push/pop operations is particularly appropriate because it is transparent to the nodes in the virtual private network and it preserves other wavelength-trunking information specific to the virtual private network.

10 Referring to Figure 19, there is shown a wavelength-switched WDM system 180 that employs both the label-push operation and the label-pop operation to provide transparent wavelength-trunking between two metropolitan area networks. The system 180 comprises a first

15 metropolitan area network 182 and a second metropolitan area network 184 that are both connected to a common long distance core network 186. The first metropolitan area network 182 includes nodes A, B, and C, while the second metropolitan area network 184 includes nodes H, I, and J.

20 The core network 186 includes a plurality of optical cross-connects (OXCs) 188. Both the first metropolitan

area network 182 and the second metropolitan area network 184 have a wavelength granularity that is sixteen times coarser than that supported by the core network 186.

It should be assumed that all the networks 182, 184, 5 and 186 share the same frequency set F_0, F_1, \dots, F_{255} . However, the frequencies used in the first metropolitan area network 182 and the second metropolitan area network 184 are of the form F_{16i} , where i is an integer between 0 and 15. Also, the different networks may use any form of 10 wavelength-switching, including any of those described above (i.e., those utilizing the label-pop operation and the label-swap operation).

In the core network 186, a virtual path (i.e., VP in Figure 19) is established from OXC D to OXC G through OXC 15 F. The virtual path carries the frequencies in the range $[F_0, F_{15}]$ between OXC D and OXC F, and in the range $[F_{64}, F_{79}]$ between OXC F and OXC G. If a label length of 4 is considered in the core network 186, the consecutive values of the topmost label of the virtual path are 0000 20 in the link between OXC D and OXC F, and 0100 in the link between OXC F and OXC G. For some connection between a

source node (S) 190 and a termination node (T) 192, a path from node B to node C to OXC D to OXC F to OXC G to node H to node J is selected. The routing and wavelength assignment algorithm allocates specific frequencies in the links of the path included in the first metropolitan area network 182 and the second metropolitan area network 184. Such assignment is based on the distribution of the load in the first metropolitan area network 182 and the second metropolitan area network 184. To transport the connection over the core network 186, the virtual path is used by selecting a frequency that meets the following conditions: 1.) It belongs to the set used by the metropolitan area networks 182 and 184 (i.e., it is of the form $F_{16,i}$); 2.) It is free over the node C to OXC D link and the OXC G to node H link; and 3.) When the first label of the virtual path is pushed into the wavelength, a wavelength is obtained that is available in the range allocated to the virtual path. It is assumed that such a frequency is F_{224} , which becomes F_{14} when label 0000 is pushed into the wavelength. In the virtual path, the allocated channel is F_{14} between OXC D and OXC F, and F_{78}

between OXC F and OXC G. At OXC G, the allocated channel is label-popped (i.e., the popped label is 0100) to obtain F_{224} . Overall, the operations in the core network 186 are completely transparent to the first metropolitan area network 182 and to the second metropolitan area network 184.

Figure 20 illustrates the two-way reservation process that takes place prior to the establishment of the connection between the source node (S) 190 and the termination node (T) 192 in the system 180 of Figure 19.

Referring to Figure 21, there is shown an optical burst-switched WDM system 200 that employs the label-pop operation in accordance with the present invention. The optical burst-switched WDM system 200 comprises a source 202, a destination 204, and a network 206, which includes a first edge node 208, a second edge node 210, and a plurality of optical cross-connects (OXC) 212 interconnected by a plurality of interconnects 214. Similar to the wavelength-switched WDM system 150 of Figure 15, in the optical burst-switched WDM system 200 of Figure 21 a number $W = 256$ (i.e., $N = 8$) of

consecutive regularly spaced wavelength channels are assumed, as is a label length $L = 2$. Also, it is assumed that the traffic demand comprises one wavelength channel between the source 202 and the destination 204. However,

5 in the system 200 of Figure 21, there is no need to establish physical cross-connections on a per-burst basis. That is, the system 200 of Figure 21 may support optical burst-switching (OBS) provided that the network nodes 212 have the two following capabilities: 1.)

10 Interpretation of OBS control messages to resolve contention among optical bursts; and 2.) Assignment of absolute priority of wavelength-switched traffic over burst-switched traffic. If such is the case, the first edge node 208 may encode multi-hop route information into

15 a burst by assigning it to the correct ingress wavelength. In networks of moderate size, all the route information may be encoded into the wavelength and decoded at network nodes 212 by the label-pop operation, as shown in Figure 21. In the example shown in Figure

20 21, at each network hop (except the last one leading to the second edge node 210), the wavelength of the optical

burst is popped according to the label-pop operation described above. Thus, the label-pop operation and the waveband routing process provide high-performance burst forwarding by eliminating the need to establish physical cross-connections on a per-burst basis. Further, on an H hop path, the minimum burst transmission latency reduces from $H(\delta_x + \delta_m) - \delta_p$ (which is the case for conventional optical burst switching systems, as described above) to $\delta_{\min} = H\delta_m - \delta_p$.

At this point it should be noted that a proper contention resolution mechanism is critical to successful transmissions. To assign absolute priority to wavelength-switched lightpaths, a simple scheme is considered where optical bursts are dropped if they are carried on wavelengths already assigned to ongoing circuits, and where burst collisions are allowed in the core network. In such a case, a burst is successfully transmitted when it does not use any busy wavelength on its path, and does not collide with other bursts. As described above, Figure 21 illustrates a successfully transmitted burst. In contrast, Figure 22 illustrates an

unsuccessful burst transmission in the system 200 of

Figure 21 due to a wavelength already having been
assigned to an ongoing circuit.

As with the system 150 of Figure 15 and the system
5 170 of Figure 17, it is easy to check that the wavelength
carrying the optical burst in the system 200 of Figures
21 and 22 is changed according to the label-pop
operation, and that the ingress wavelength carries the
whole route information. More sophisticated contention
10 resolution schemes may be considered within the scope of
the present invention where contending optical bursts are
buffered or deflected.

As described above, the label-pop operation enables
the efficient implementation of source routing for
15 optical burst-switching (OBS). However, labels may also
correspond to virtual paths that are themselves label-
switched. When the topmost label of an optical burst
maps to such a virtual path, the label stack of the
virtual path may be pushed (i.e., via the label-push
20 operation) into the stack of the burst to enable proper
forwarding. Thus, optical burst-switching (OBS) may also

be efficiently implemented utilizing the label-push operation in accordance with the present invention.

In view of the foregoing, the new label-switching paradigm described herein supports wavelength-switching, wavelength-trunking, as well as optical burst-switching, by encoding part or all of a network route into the wavelength assigned to bursts or circuits at a network ingress node. The new label-switching approach is a new development in the area of waveband routing where, at network nodes, disjoint wavebands are assigned to outgoing interfaces, and incoming channels are routed to outgoing interfaces according to their incoming interfaces and to the band to which they belong. Typically, the wavebands form a partition of the available spectrum. When optical frequencies are mapped to corresponding binary numbers, the waveband of a particular optical frequency is identified by a specific bit pattern, which corresponds to a routing label.

Waveband routing is quite appropriate for optical burst-switching (OBS) because it does not require physical cross-connection on a per-burst basis in optical

cross-connects (OXCs). However, when using waveband routing, a major problem is the composition of the end-to-end routing service, when the bands are defined locally. In the worst-case, the different wavebands needed along specific paths may not intersect, therefore preventing the use of waveband routing for the selected paths.

An important difference between the present invention and previous waveband routing schemes lies in a solution to the waveband-intersection problem. The solution is based on a systematic transformation of routing labels at each network node to enable proper forwarding at the next-hop. These transformations use specific operations, which we call label-swap, label-pop, and label-push to be consistent with the MPLS model (see D. Awduche, "MPLS traffic engineering in IP networks", IEEE Communications magazine, vol. 37, pages 42-47, Dec. 1999).

The label-swap operation involves replacing a current routing label by another routing label that was previously advertised by the signaling protocol at

connection setup.

The label-pop operation occurs when the binary number corresponding to a given optical frequency is partitioned into several consecutive bit-fields, which correspond to routing labels, and form a label stack. The label-pop operation involves popping the label stack by removing the topmost label, and by shifting the position of all of the other labels of the stack by a number of bits equal to the size of the removed label.

10 The advantage of the label-pop operation is to enable source routing by encoding multiple routing labels when proper wavelengths are selected at a network ingress node.

The label-push operation inserts new routing labels into the stack. In combination, label-pop and label-push operations also provide good means to support wavelength-trunking, when network ingress nodes with classifying functions push labels into wavelengths according to virtual lightpaths through which they must be routed.

20 All the above-mentioned label-stack operations only require wavelength-conversion devices. For these

reasons, they may be implemented by all-optical means,
which provide high data packet forwarding performance.
The new label-switching paradigm decouples the optical
data forwarding process from the wavelength-contention

5 problem. Therefore, network nodes are able to route
optical packets or bursts on demand, independent of
quality-of-service (QoS) levels that result from the
contention resolution process. Despite this added
flexibility, the new label-switching paradigm can still
10 support traditional wavelength-switching. Overall, the
new label-switching paradigm enables the design of
networks possessing mixed wavelength-switching and
optical burst-switching (OBS) capabilities supported on
the same network infrastructure.

15 The present invention is not to be limited in scope
by the specific embodiments described herein. Indeed,
various modifications of the present invention, in
addition to those described herein, will be apparent to
those of ordinary skill in the art from the foregoing
20 description and accompanying drawings. Thus, such
modifications are intended to fall within the scope of

10

CLAIMS

What is claimed is:

1 1. A method for routing data within an optical network
2 having a plurality of network nodes, the method
3 comprising the steps of:

4 receiving data at a first network node via a first
5 optical signal having a first wavelength, the first
6 wavelength corresponding to a first optical frequency,
7 the first optical frequency being mapped to a first
8 binary representation, the first binary representation
9 being divided into a first plurality of fields, at least
10 one of the first plurality of fields corresponding to a
11 routing label in a first label stack, a top routing label
12 in the first label stack indicating a second network
13 node; and

14 based at least partially upon the top routing label,
15 transmitting the data from the first network node to the
16 second network node via a second optical signal having a
17 second wavelength.

1 2. The method as defined in claim 1, further comprising
2 the step of:

3 popping the top routing label off the first label
4 stack so as to promote a next routing label in the first
5 label stack.

1 3. The method as defined in claim 2, wherein the second
2 wavelength corresponds to a second optical frequency, the
3 second optical frequency being mapped to a second binary
4 representation, the second binary representation being
5 divided into a second plurality of fields, at least one
6 of the second plurality of fields corresponding to a
7 routing label in a second label stack, a top routing
8 label in the second label stack indicating a third
9 network node.

1 4. The method as defined in claim 3, wherein the top
2 routing label in the second label stack corresponds to
3 the next routing label in the first label stack.

1 5. The method as defined in claim 4, wherein the
2 network accommodates 2^N frequencies in the form of,

3
$$f_i = f_0 + i \cdot \Delta f$$

4 wherein $i = 0, 1, \dots, 2^N - 1$, wherein the second optical
5 frequency is defined by,

6
$$f_{i_{out}} = f_0 + 2^L \left((f_{i_{in}} - f_0) - 2^{N-L} L \Delta f \right)$$

7 and,

8
$$i_{out} = 2^L \left(i_{in} - 2^{N-L} L \right)$$

9 wherein $f_{i_{in}}$ represents the first optical frequency, i
10 represents the value of the top routing label in the
11 first label stack, and L represents the bit length of the
12 top routing label in the first label stack.

1 6. The method as defined in claim 1, further comprising
2 the step of:

3 swapping the top routing label in the first label
4 stack with a new routing label when the first label stack
5 contains more than two routing labels.

1 7. The method as defined in claim 6, wherein the second
2 wavelength corresponds to a second optical frequency, the

3 second optical frequency being mapped to a second binary
4 representation, the second binary representation being
5 divided into a second plurality of fields, at least one
6 of the second plurality of fields corresponding to a
7 routing label in a second label stack, a top routing
8 label in the second label stack indicating a third
9 network node.

1 8. The method as defined in claim 7, wherein the top
2 routing label in the second label stack corresponds to
3 the new routing label.

1 9. The method as defined in claim 8, wherein the
2 network accommodates 2^N frequencies in the form of,

3
$$f_i = f_0 + i \cdot \Delta f$$

4 wherein $i = 0, 1, \dots, 2^N - 1$, wherein the second optical
5 frequency is defined by,

6
$$f_{i_{out}} = f_{i_{in}} + 2^{N-L} \left(-l^1 \right) \Delta f$$

7 and,

8
$$i_{out} = i_{in} + 2^{N-L} \left(-l^1 \right)$$

9 wherein $f_{i_{in}}$ represents the first optical frequency, l^1

10 represents the value of the top routing label in the
11 first label stack, l represents the value of the new
12 routing label, and L represents the bit length of the top
13 routing label in the first label stack.

1 10. The method as defined in claim 1, further comprising
2 the step of:
3 pushing a new routing label onto the first label
4 stack.

1 11. The method as defined in claim 10, wherein the
2 second wavelength corresponds to a second optical
3 frequency, the second optical frequency being mapped to a
4 second binary representation, the second binary
5 representation being divided into a second plurality of
6 fields, at least one of the second plurality of fields
7 corresponding to a routing label in a second label stack,
8 a top routing label in the second label stack indicating
9 a third network node.

1 12. The method as defined in claim 11, wherein the top
2 routing label in the second label stack corresponds to
3 the new routing label.

1 13. The method as defined in claim 12, wherein the
2 network accommodates 2^N frequencies in the form of,

3
$$f_i = f_0 + i \cdot \Delta f$$

4 wherein $i = 0, 1, \dots, 2^N - 1$, wherein the second optical
5 frequency is defined by,

6
$$f_{i_{out}} = f_0 + \left\lfloor \frac{(i_{in} - f_0)}{\Delta f} \right\rfloor 2^{-L} \Delta f + 2^{N-L} \Delta f$$

7 and,

8
$$i_{out} = \left\lfloor \frac{i_{in}}{2^L} \right\rfloor + 2^{N-L}$$

9 wherein $f_{i_{in}}$ represents the first optical frequency, i
10 represents the value of the top routing label in the
11 second label stack, and L represents the bit length of
12 the top routing label in the second label stack.

1 14. The method as defined in claim 1, wherein the first
2 wavelength is the different from the second wavelength.

1 15. The method as defined in claim 1, wherein the first
2 wavelength is the same as the second wavelength.

1 16. The method as defined in claim 1, wherein at least
2 another one of the first plurality of fields corresponds
3 to a termination field indicating an end of the first
4 label stack.

1 17. The method as defined in claim 1, wherein at least
2 another one of the first plurality of fields corresponds
3 to a contention field for differentiating the first
4 wavelength from a third wavelength.

1 18. The method as defined in claim 17, wherein the data
2 is a first data, wherein second data is received at the
3 first network node via a third optical signal having the
4 third wavelength, and wherein the first optical signal
5 and the third optical signal have similar routing paths
6 through the network.

1 19. An apparatus for routing data within an optical
2 network having a plurality of network nodes, the
3 apparatus comprising:

4 an optical receiver for receiving data at a first
5 network node via a first optical signal having a first
6 wavelength, the first wavelength corresponding to a first
7 optical frequency, the first optical frequency being
8 mapped to a first binary representation, the first binary
9 representation being divided into a first plurality of
10 fields, at least one of the first plurality of fields
11 corresponding to a routing label in a first label stack,
12 a top routing label in the first label stack indicating a
13 second network node; and

14 an optical transmitter for transmitting, based at
15 least partially upon the top routing label, the data from
16 the first network node to the second network node via a
17 second optical signal having a second wavelength.

1 20. The apparatus as defined in claim 19, wherein the
2 first wavelength is the different from the second
3 wavelength.

1 21. The apparatus as defined in claim 19, wherein the
2 first wavelength is the same as the second wavelength.

1 22. The apparatus as defined in claim 19, wherein at
2 least another one of the first plurality of fields
3 corresponds to a termination field indicating an end of
4 the first label stack.

1 23. The apparatus as defined in claim 19, wherein at
2 least another one of the first plurality of fields
3 corresponds to a contention field for differentiating the
4 first wavelength from a third wavelength.

1 24. The apparatus as defined in claim 23, wherein the
2 data is a first data, wherein second data is received at
3 the first network node via a third optical signal having
4 the third wavelength, and wherein the first optical
5 signal and the third optical signal have similar routing
6 paths through the network.

ABSTRACT OF THE DISCLOSURE

A technique for routing data within an optical network having a plurality of network nodes is disclosed.

In one embodiment, the technique is realized by
5 receiving data at a first network node via a first
optical signal having a first wavelength. The first
wavelength corresponds to a first optical frequency, and
the first optical frequency is mapped to a first binary
representation. The first binary representation is
10 divided into a first plurality of fields, wherein at
least one of the first plurality of fields corresponds to
a routing label in a first label stack. A top routing
label in the first label stack indicates a second network
node. Based at least partially upon the top routing
15 label, the data is transmitted from the first network
node to the second network node via a second optical
signal having a second wavelength. The first wavelength
may be either the same as or different from the second
wavelength.

30

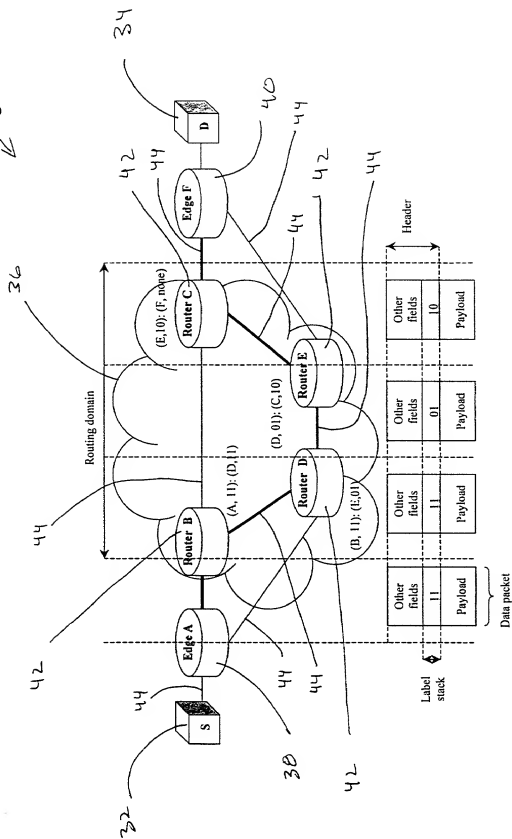


Figure 5 (prior art)

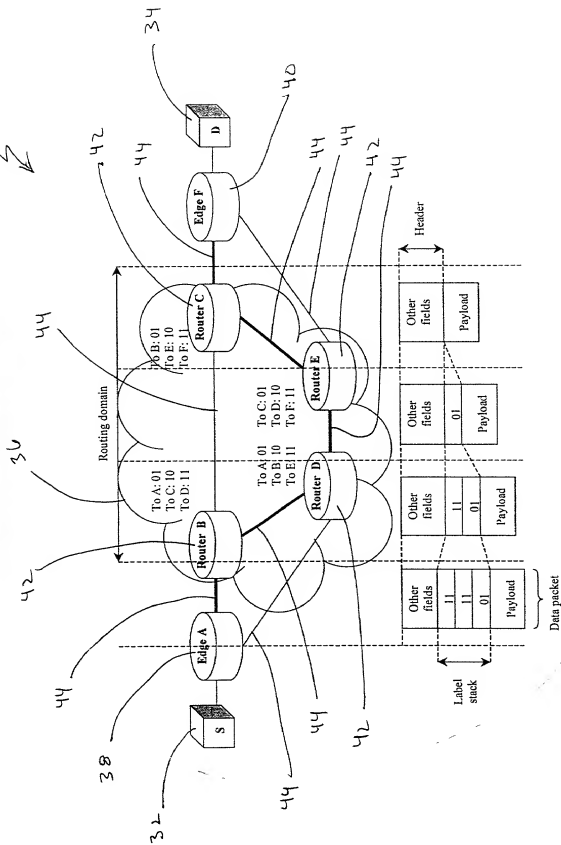


Figure 6 (prior art)

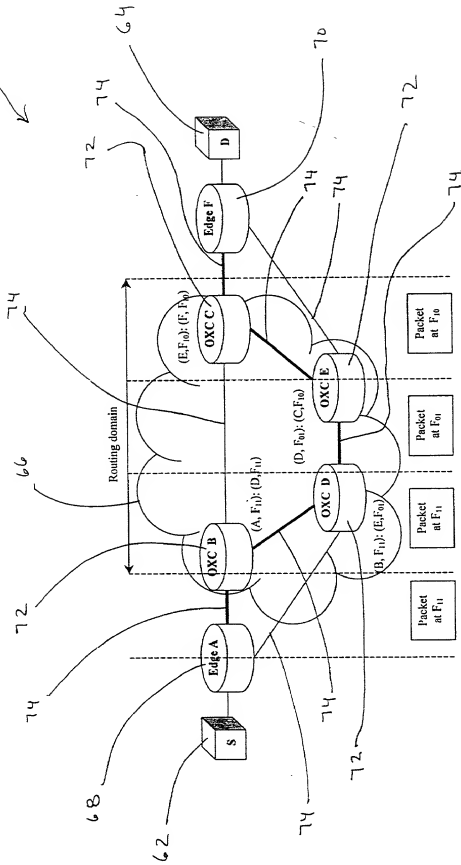


Figure 7 (prior art)

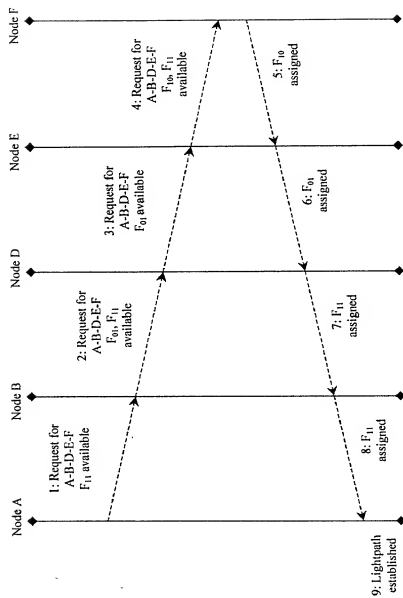


Figure 8 (prior art)



Figure 9 (prior art)

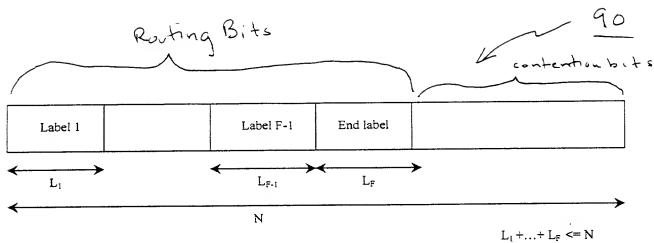


Figure 10

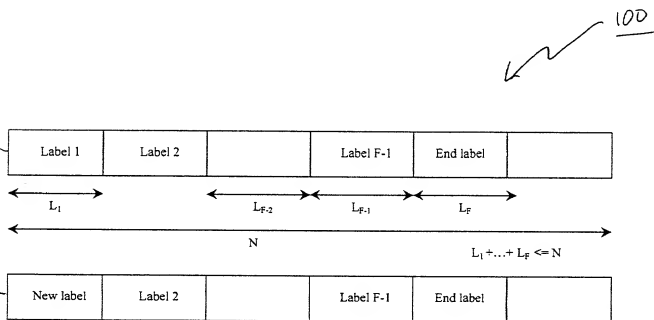


Figure 11

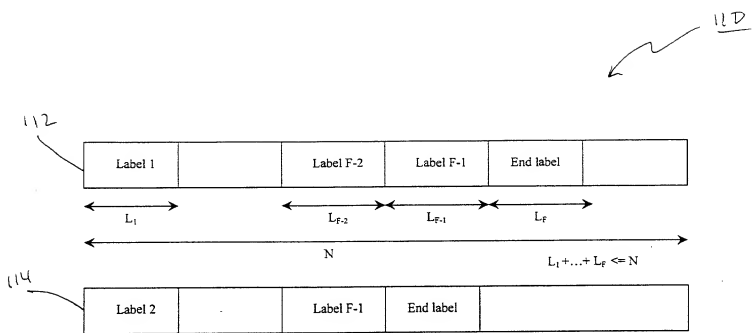


Figure 12

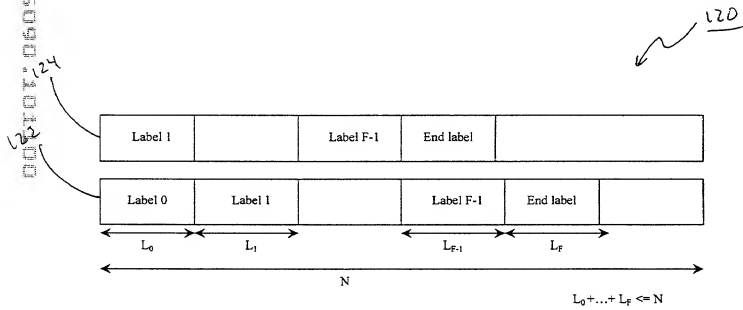


Figure 13

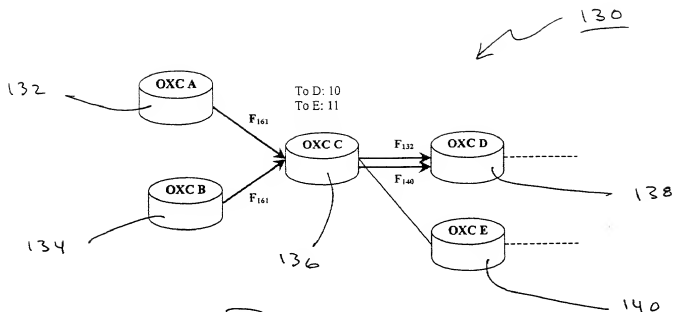


Figure 14

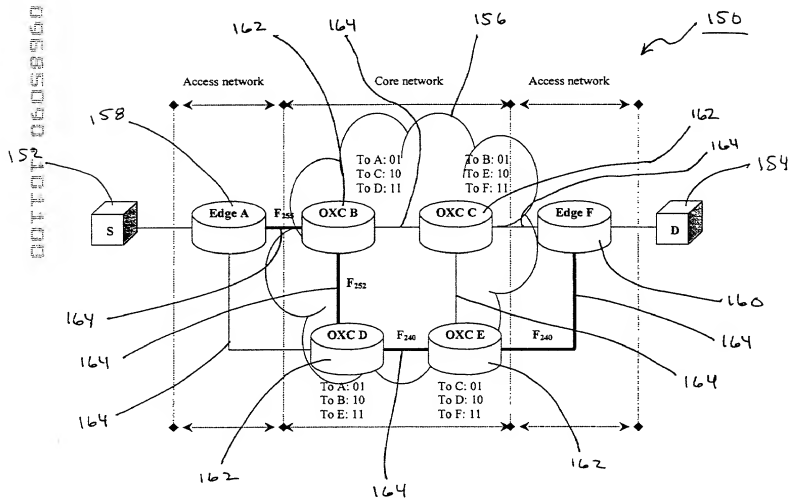


Figure 15

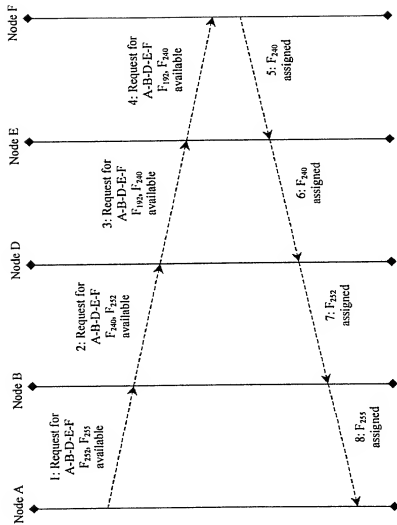


Figure 16

170

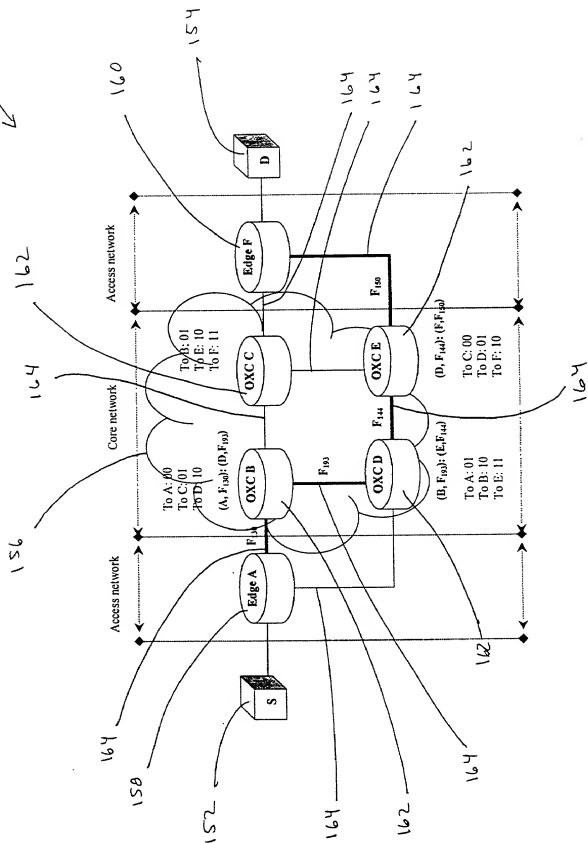


Figure 17

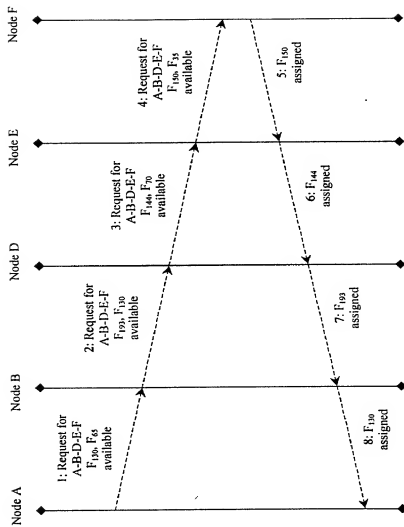


Figure 18

180 ~~~~~

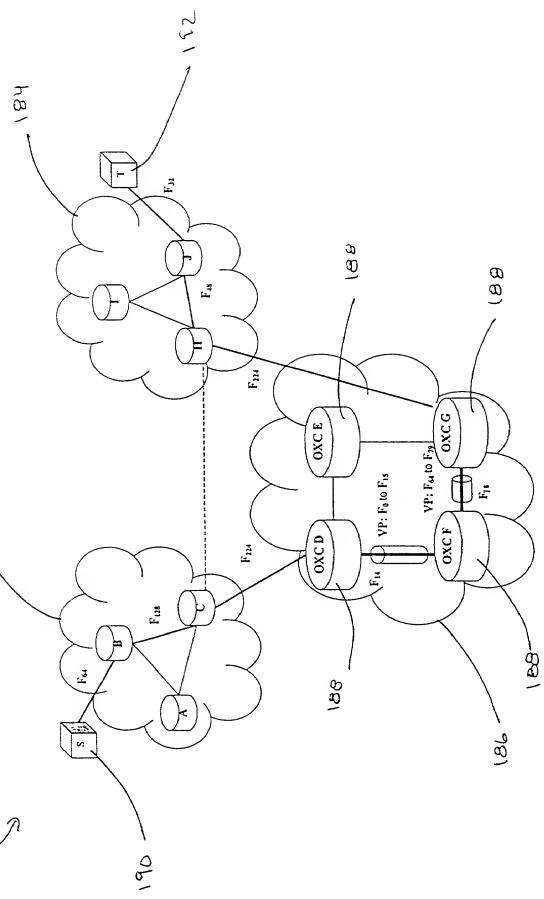


Figure 19

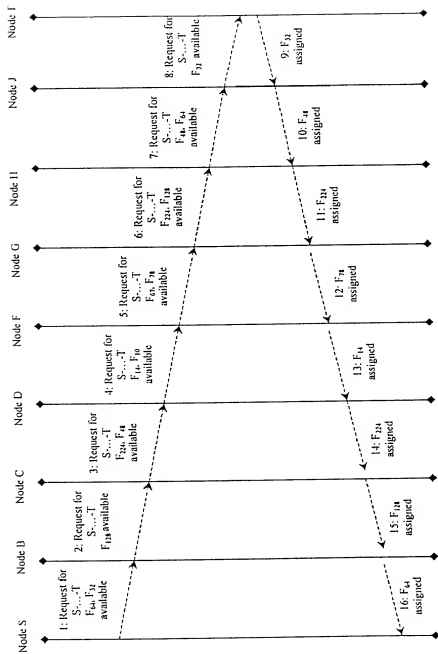


Figure 20

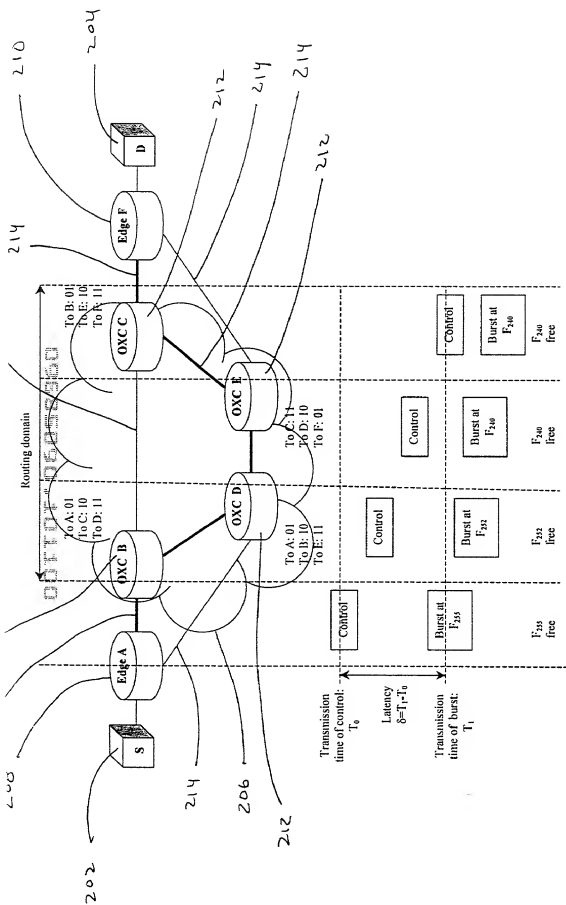


Figure 21

Figure 22

COMBINED DECLARATION AND POWER OF ATTORNEY

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name; and

I believe that I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: "Technique for Routing Data Within an Optical Network", the specification of which:

(X only one) (a) x is attached hereto.
(b) _____ was filed as U.S. Patent Application No. _____ on _____, 20____, and was amended on _____, 20____ (if applicable).

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claims as amended by any amendment referred to above.

I hereby acknowledge the duty to disclose to the Office all information known to me to be material to the patentability of this application as defined in 37 CFR § 1.56.

I hereby claim foreign priority benefits under 35 U.S.C. § 119 of any foreign application(s) for patent or inventor's certificate listed below, and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of any application on which priority is claimed:

Application No.	Country	Filing Date (MM/DD/YY)	Priority Claimed
_____	_____	_____	_____
_____	_____	_____	_____

I hereby claim the benefit under 35 U.S.C. § 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not

Patent Application
Attorney Docket No.: 57983.000005
Client Reference No.: 12283RO

disclosed in the prior United States application in the manner provided by the first paragraph of 35 U.S.C. § 112, I hereby acknowledge the duty to disclose to the Office all information known to me to be material to the patentability of this application as defined in 37 CFR § 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

Application No.	Filing Date (MM/DD/YY)	Status (Patented, Pending, Abandoned)
-----------------	---------------------------	--

_____	_____	_____
_____	_____	_____

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

I hereby appoint: Thomas E. Anderson, Reg. No. 37,063; Thomas J. Scott, Jr., Registration No. 27,836; Stanislaus Aksman, Registration No. 28,562; James G. Gatto, Registration No. 32,694; Christopher C. Campbell, Registration No. 37,291; Thomas M. Blasey, Registration No. 33, 475; Henry C. Su, Registration No. 37,738; Brian M. Buroker, Registration No. 39,125; Charles F. Hollis, Registration No. 40,650; Jonathan D. Link, Registration No. 41,548; Kevin J. Dunleavy, Registration No. 32,024; Kevin T. Duncan, Registration No. 41,495; Charles B. Lobsenz, Registration No. 37,857; George Georgellis, Registration No. 43,632; Stephen T. Schreiner, Registration No. 43,097; Christopher J. Cuneo, Registration No. 42,450; Raphael A. Valencia, Registration No. 43,216; Scott D. Balderston, Registration No. 35,436; Steven P. Klocinski, Registration No. 39,251; Yisun Song, Registration No. 44,487; Jennifer A. Albert, Registration No. 32,012; Kerry Owens, Registration No. 37,412; Milan M. Vinnola, Registration Number 45,979; Devin S. Morgan, Registration No. 45,562; Andrew J. Ririe, Registration No. P45,979; Carl L. Benson, Registration No. 38,378; Robin C. Clark, Registration No. 49,050; Herbert V.

05685090-10110

Patent Application
Attorney Docket No.: 57983.000005
Client Reference No.: 12283RO

Kerner, Registration No. 42,721; Rene' Vazquez, Registration No. 38,647, all of the firm **Hunton & Williams**, 1900 K Street, NW, Washington, DC 20006, as my attorneys and/or agents, with full power of substitution and revocation, to prosecute this application, and any provisionals, continuations, continuations-in-part, divisionals, appeals, reissues, substitutions, re-examinations, and extensions thereof, and to transact all business in the United States Patent and Trademark Office connected therewith, to appoint any individuals under an associate power of attorney, and to file and prosecute any international patent application based thereon before any international authorities.


Please address all correspondence and direct all telephone calls to:

Thomas E. Anderson
Hunton & Williams
1900 K Street, NW
Washington, DC 20006
(202) 955-1808
(202) 778-2201 (Fax)

NAMED INVENTOR(S) :

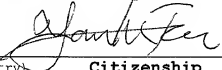
0630-0101

1

2	Full Name Delfin Y. Montuno	Inventor's Signature 	Date Oct. 5, 2000
	Residence (city, state, country) Citizenship 122 Knudson Drive Kanata, ON K2K2L8 Canada Canada		
	Post Office Address (include zip code) 122 Knudson Drive Kanata, ON K2K2L8 Canada		

001101-06058960

Patent Application
Attorney Docket No.: 57983.000005
Client Reference No.: 12283RO

3	Full Name Yanhe Fan	Inventor's Signature 	Date Sept. 28, 2000
	Residence (city, state, country) 25 Woodridge Cr. Apt. 1212 Neapen, ON K2B7T4 Canada		Citizenship Canada
	Post Office Address (include zip code) 25 Woodridge Cr. Apt. 1212 Neapen, ON K2B7T4 Canada		

00000000-101100